
Subject: Re: [PATCH][for -mm] Fix and Enhancements for memory cgroup [3/6] add helper function for page_cgroup

Posted by [KAMEZAWA Hiroyuki](#) on Wed, 10 Oct 2007 00:38:39 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Tue, 9 Oct 2007 20:26:42 +0900

KAMEZAWA Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com> wrote:

```
> >
> > > + */
> > > + if (clear_page_cgroup(page, pc) == pc) {
> >
> > OK.. so we've come so far and seen that pc has changed underneath us,
> > what do we do with this pc?
> >
> > Hmm... How about this ?
> ==
> if (clear_page_cgroup(page, pc) == pc) {
> /* do usual work */
> } else {
> BUG();
> }
> == or BUG_ON(clear_page_cgroup(page, pc) != pc)
>
> I have no clear idea when this race will occur.
```

After good sleep, I noticed there is a race with force_reclaim (in patch 6).

force_reclaim doesn't check refcnt before clearing page->pc.

My final view will be

```
==
if (clear_page_cgroup(page, pc) == pc) {
/* do usual work */
} else {
/* force reclaim clears page->page_cgroup */
}
```

==
Anyway, I'll add a meaningful comment here.

Thanks,

-Kame

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
